

CASCADED POLICING SYSTEMS AND METHODS

Field of the Invention

The invention relates to the policing of data flows, for example flows of IP (Internet Protocol) packets, in a manner delivering class of service, also referred to as quality of service, differentiability.

Background of the Invention

It is now a common objective in IP networks to provide the option of a guaranteed QoS (quality of service). See for example, 1) "Quality of service in ATM networks : State-of-the-art Traffic Management", Natalie Giroux, Sudhakar Ganti, 1999 by Prentice-Hall PTR, pages 38-46 and 61; 2) "Specification of Guaranteed Quality of Service", Shenker, et al, RFC 2212, Standards Track, September 1997, pages 1 to 20; and 3) "An Architecture for Differentiated Services", Blake, et al, RFC 2475, Informational, December 1998, pages 1 to 36.

Associated with QoS delivery is the concept of traffic "policing", (also synonymous with "marking" or "metering") whereby a service provider ensures that at the same time a customer is receiving the QoS paid for, they are not in certain respects exceeding that QoS.

Referring now to Figure 1, shown is an example of a customer's traffic source 10 generating traffic 14 which is sent to a network 12 through a connection 15. During the setup of such a connection 15, typically the customer has requested/negotiated certain traffic parameters for the traffic 14, such as bandwidth, delay etc., and pays for the connection accordingly. The network 12 has a policing node 16 at which the traffic 14 is policed in accordance with the negotiated

parameters. Typically the policing node 16 is the first point of access within the network 12 for the traffic 14.

The policing node 16 has a policer (synonymous with "meter" or "marker") 18 responsible for marking packets which constitute traffic 14 as either conforming, non-conforming, or partially conforming. The policer 18 is typically implemented using a leaky bucket mechanism. Each time a packet of traffic 14 arrives, a bucket is filled by a number of policing units, or tokens, corresponding to an allowed burst of data. The bucket continuously leaks tokens at a rate reflective of the bandwidth or rate to be provided. In the event the bucket overflows, packets are marked as non-conforming. Packets which arrive while the bucket is not overflowing are marked as conforming. Typically, allowances are made by the policer 18 to realize both an average rate (sometimes referred to as the committed information rate or CIR), and a burst tolerance (BT). Burst tolerance can be provided for example by allowing the bucket to accumulate up to the maximum token bucket size. This allows packets to be transmitted at a rate greater than the average for a short period of time.

Existing policing algorithms are designed to police a single traffic flow to a single set of negotiated specifications. When there are multiple traffic flows from a single customer, multiple independent policers have been employed.

Summary of the Invention

Embodiments of the invention provide cascaded policing methods and systems which allow lower priority traffic to benefit from otherwise unused capacity allocated to higher priority traffic of a given customer/service with multiple classes of service.

A first broad aspect of the invention provides a method of policing packet traffic. The method involves policing packets of a first class in accordance with at least one policing parameter associated with the first class, and
5 policing packets of a second class in accordance with at least one policing parameter associated with the second class in a manner which gives to the second class at least a portion of a traffic throughput afforded to the first class by at least one of said at least one policing parameter associated with the
10 first class of traffic which is not being used by the packets of the first class.

The policing parameters under consideration might for example be rate guarantees provided to different traffic classes. The policing parameters might also include burst
15 tolerances of the different traffic classes.

The method is easily adapted to an arbitrary number of different traffic classes.

Another broad aspect of the invention provides a method of policing traffic involving defining a traffic class
20 rate guarantee for each of a plurality of traffic classes to be provided by a service, and a service rate guarantee for the service, and policing combined traffic containing traffic of each of the plurality of traffic classes in a manner which guarantees each class its respective traffic class rate
25 guarantee, and in a manner which guarantees the service rate guarantee for the combined traffic. This effectively amounts to a two-tier rate guarantee.

Preferably each of a respective combined traffic comprising a given traffic class plus all conforming higher
30 class traffic, the policing being done at a rate equal to the traffic class rate guarantee for that traffic class plus the

78945-7-435350

traffic class rate guarantees for at least one and preferably all higher classes of traffic.

In one embodiment, a method of policing a plurality N of traffic classes C_i , each having a respective rate guarantee R_i , $i=1, \dots, N$, $N \geq 2$ is provided. The method involves policing traffic of class C_1 according to rate R_1 , and for each other class C_i , policing traffic of class C_i plus conforming traffic of class(es) C_1, \dots, C_{i-1} according to an aggregate rate $RA_i =$

$$\sum_{i=1}^N R_i.$$

This method may be adapted to include consideration of burst tolerance. For example if each traffic class C_i has a respective burst tolerance BT_i , the method preferably further involves policing traffic of class C_1 according to BT_1 , and for each other class C_i , policing traffic of class C_i plus conforming traffic of class(es) C_1, \dots, C_{i-1} according to an aggregate burst tolerance $BA_i = \sum_{i=1}^N BT_i$.

Embodiments of the invention also provide a policer which might be any suitable combination of hardware and/or software, and a network node adapted to implement any of the above described methods. A processing platform readable medium having stored thereon instructions for a processing platform to implement any of the above described methods is also provided.

Brief Description of the Drawings

Preferred embodiments of the invention will now be described with reference to the attached drawings in which:

Figure 1 is a schematic diagram of a conventional policing arrangement;

Figure 2 is a schematic diagram of a system in which traffic is policed according to a method provided by an embodiment of the invention;

Figure 3 is a logical view of the functionality of the cascaded policer of Figure 2; and

Figure 4 is a traffic flow diagram illustrating a preferred method of implementing the cascaded policer of Figure 2.

Detailed Description of the Preferred Embodiments

Embodiments of the invention provide for the aggregate policing of multiple traffic classes within a service. A service is defined as a data communications path through a network. It is desirable to provide class of service differentiation within a service. Class of service differentiation involves treating sub-flows of packets generated within the service in a different manner. Referring now to Figure 2, shown is an example of a traffic source associated with service 24. A customer subscribing to the service 24 provided by network 22 generates traffic at traffic source 20. The service 24 includes four traffic classes, indicated logically by class C1 traffic 26, class C2 traffic 28, class C3 traffic 30 and class C4 traffic 32 flowing between the traffic source 20 and the network 22. The traffic classes 26,28,30,32 collectively constitute the service 24 being provided. Although Figure 1 only shows traffic ingress to the network 22, complete service delivery would involve delivering the traffic through the network to one or more destinations.

In a preferred embodiment of the invention, the traffic consists of IP packets, and the traffic classes might for example be IETF (Internet Engineering Task Force) DiffServe (Differentiated Services) classes EF (expedited forwarding),

AF1 (assured forwarding 1), AF2 (assured forwarding 2), and BE (best effort). Of course, other packet types and traffic classes may alternatively be employed, such as ATM and Frame Relay.

5 During the setup of such a service 24, certain traffic parameters are requested/negotiated for each of the traffic classes, such as bandwidth, delay etc., and the service is paid for accordingly. The network 22 has a policing node 34 at which the traffic associated with each traffic class service
10 24 is policed in accordance with the negotiated parameters. Typically the policing node 34 is the first point of access within the network 22 for the traffic of service 24. Policing of the traffic classes 26,28,30,32 within policing node 34 is performed by a cascaded policer 38 which outputs marked traffic
15 39.

 A logical view of the functionality of the cascaded policer 38 is provided in Figure 3. The four traffic classes 26,28,30,32 are shown entering the cascaded policer 38. According to this embodiment of the invention, policing is
20 performed by the cascaded policer 38 in a manner such that if a higher priority class does not use the full capacity rate allocated (and thus paid for), unused capacity is allowed to be used by lower classes. For the purpose of this example, it is assumed that the order of priority for the traffic classes from
25 highest to lowest is Class C1, Class C2, Class C3 and then Class C4. It is assumed that for Class C1, a CIR of R1 has been paid for, meaning that regardless of what is going on with the other classes, Class C1 is going to be allowed to transmit R1. Similarly, it is assumed that for Class C2, a CIR of R2
30 has been paid for, meaning that regardless of what is going on with the other classes, Class C2 is going to be allowed to transmit R2. It is assumed that for Class C3, a CIR of R3 has been paid for, meaning that regardless of what is going on with

the other classes, Class C3 is going to be allowed to transmit R3. It is assumed that Class C4 is a best effort class which has a guaranteed CIR of R4 (which may be zero).

The policing is to be performed in accordance with
5 the following rules:

Class C1 traffic < R1;

Conforming Class C1 + Class C2 < R1 + R2;

Conforming Class C1 + Conforming Class C2 + Class C3
< R1 + R2 + R3;

10 Conforming Class C1 + Conforming Class C2 +
Conforming Class C3 + Class C4 < R1 + R2 + R3 + R4.

Another way of expressing this for an arbitrary
number N of classes is as follows:

police traffic of class C1 according to rate R1;

15 for each other class Ci, police traffic of class Ci
plus conforming traffic of class(es) C1,...,Ci-1 according to an
aggregate rate $RA_i = \sum_{i=1}^N R_i$.

In the above, the first rule means that class C1
traffic is policed to R1. Traffic beyond R1 will be marked as
20 non-conforming. Traffic below R1 will be marked as conforming.

The second rule effectively means that class C2
traffic is policed to R1 + R2 - conforming class C1 traffic.
Traffic beyond this amount will be marked as non-conforming.
Traffic below this amount will be marked as conforming.

25 The third rule effectively means that class C3
traffic is policed to R1 + R2 + R3 - conforming class C1

traffic - conforming class C2 traffic. Traffic beyond this amount will be marked as non-conforming. Traffic below this amount will be marked as conforming.

Finally, the fourth rule effectively means that class
5 C4 traffic is policed to $R1 + R2 + R3 + R4$ - conforming class 1
traffic - conforming class C2 traffic - conforming class C3
traffic. Traffic beyond this amount will be marked as non-
conforming. Traffic below this amount will be marked as
conforming.

10 The effect of policing in this manner is that a
customer has paid for an amount $R1$ of class C1 traffic
capacity, and if this is not used, rather than policing class
C2 at its nominal rate of $R2$, class C2 traffic is given the
opportunity to be transmitted on the left over capacity paid
15 for class C1 and so on.

Effectively, a two-tier rate guarantee mechanism is
provided, with each class of service being given its own
respective rate guarantee, and the service as a whole also
being given a rate guarantee which is equal to the sum of the
20 individual rate guarantees.

There are many ways of practically achieving these
rules. One example is given in the traffic flow diagram of
Figure 4. Class C1 traffic 24 enters a first policer 50 which
marks traffic as either conforming or non-conforming according
25 to rate $R1$. The non-conforming traffic may be dropped right
there, or may be left in the packet stream for the network to
decide what to do with it at a later time. The traffic thus
marked 52, and class C2 traffic 26 enters a second policer 54
which polices the combination of class C2 traffic 26 and
30 conforming class C1 traffic at $R1 + R2$. Any non-conforming
class C1 traffic in marked traffic 52 is ignored. Conforming
Class C1 traffic is already marked as conforming, so only class

C2 traffic can be marked non-conforming by the second policer 54 producing marked traffic 56. Then, the combination of conforming class C1 and conforming class C2 and class C3 traffic 28 is policed at R1 + R2 + R3 by a third policer 58 5 producing marked traffic 60. Finally, the combination of conforming class C1, conforming class C2, conforming class C3, and class C4 traffic 30 is policed at R1 + R2 + R3 + R4 by a fourth policer 62 producing marked traffic 39.

Preferably, the burst tolerance is cascaded in the 10 same manner as the committed information rates. Thus, if in the absence of any other considerations class C1, class C2, class C3 and class C4 have burst tolerances of BT1, BT2, BT3, and BT4 respectively, then the policing is performed such that class C1 is given a burst tolerance of BT1, the combination of 15 class C1 and class C2 is given a burst tolerance of BT1 + BT2, the combination of class C1, class C2 and class C3 is given a burst tolerance of BT1 + BT2 + BT3, and finally, the combination of class C1, class C2, class C3 and class C4 is given a burst tolerance of BT1 + BT2 + BT3 + BT4.

20 Mathematically, this can be expressed as follows for an arbitrary number N of traffic classes:

police traffic of class C1 according to BT1;

for each other class Ci, policing traffic of class Ci plus conforming traffic of class(es) C1,...,Ci-1 according to an

25 aggregate burst tolerance $BA_i = \sum_{i=1}^N BT_i$.

Specific examples have been given in which both the committed information rate and the burst tolerance of multiple traffic classes are considered in an aggregate manner. There may be other parameters which may be similarly cascaded.

In the described embodiment, there are four traffic classes which are being policed by the cascaded policer. More generally, any number of traffic classes may be policed in this manner.

5 Also, the above described embodiment, specific mechanisms and methods of allocating all of a class's unused capacity to lower priority classes have been provided. More generally, embodiments of the invention include any method of policing which results in some or all of a class's unused
10 capacity being made available to lower priority classes.

Numerous modifications and variations of the present invention are possible in light of the above teachings. It is therefore to be understood that within the scope of the appended claims, the invention may be practised otherwise than
15 as specifically described herein.

78945-7